ACM PORTAL
USPTO
Subscribe (Full Service) Register (Limited Service, Free) Login
Search: The ACM Digital Library The Guide
+arbitrating +interconnect +design +router +target +languag SEARCH
THE ACM DIGITAL LIBRARY
Feedback Report a problem Satisfaction survey

Terms used
arbitrating interconnect design router target language hardware

Found **52** of **158,639**

Sort results by relevance
Display results expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 52 Result page: **1** 2 3 next

Relevance scale

1 A comparative study of arbitration algorithms for the Alpha 21364 pipelined router
Shubhendu S. Mukherjee, Federico Silla, Peter Bannon, Joel Emer, Steve Lang, David Webb
October 2002 **Proceedings of the 10th international conference on Architectural support for programming languages and operating systems**, Volume 30 , 36 , 37 Issue 5 , 5 , 10
Full text available: pdf(1.44 MB) Additional Information: full citation, abstract, references

Interconnection networks usually consist of a fabric of interconnected routers, which receive packets arriving at their input ports and forward them to appropriate output ports. Unfortunately, network packets moving through these routers are often delayed due to conflicting demand for resources, such as output ports or buffer space. Hence, routers typically employ *arbiters* that resolve conflicting resource demands to maximize the number of matches between packets waiting at input ports an ...

2 Special session on reconfigurable computing: The happy marriage of architecture and application in next-generation reconfigurable systems
Ingrid Verbauwhede, Patrick Schaumont
April 2004 **Proceedings of the 1st conference on Computing frontiers**
Full text available: pdf(398.28 KB) Additional Information: full citation, abstract, references, index terms

New applications and standards are first conceived only for functional correctness and without concerns for the target architecture. The next challenge is to map them onto an architecture. Embedding such applications in a portable, low-energy context is the art of molding it onto an energy-efficient target architecture combined with an energy efficient execution. With a reconfigurable architecture, this task becomes a two-way process where the architecture adapts to the application and vice-vers ...

Keywords: embedded, real-time systems

3 Special issue: AI in engineering
D. Sriram, R. Joobbani
January 1985 **ACM SIGART Bulletin**, Issue 91
Full text available: pdf(8.79 MB) Additional Information: full citation, abstract

The papers in this special issue were compiled from responses to the announcement in the July 1984 issue of the SIGART newsletter and notices posted over the ARPAnet. The interest

being shown in this area is reflected in the sixty papers received from over six countries. About half the papers were received over the computer network.

4 Balancing performance and flexibility with hardware support for network architectures
Ilija Hadžić, Jonathan M. Smith
November 2003 **ACM Transactions on Computer Systems (TOCS)**, Volume 21 Issue 4
Full text available: pdf(719.03 KB) Additional Information: full citation, abstract, references, index terms

The goals of performance and flexibility are often at odds in the design of network systems. The tension is common enough to justify an architectural solution, rather than a set of context-specific solutions. The Programmable Protocol Processing Pipeline (P4) design uses programmable hardware to selectively accelerate protocol processing functions. A set of field-programmable gate arrays (FPGAs) and an associated library of network processing modules implemented in hardware are augmented with so ...

Keywords: FPGA, P4, computer networking, flexibility, hardware, performance, programmable logic devices, programmable networks, protocol processing

5 Multiprocessor SoC: design strategies and programming models: Modeling operation and microarchitecture concurrency for communication architectures with application to retargetable simulation
Xinping Zhu, Wei Qin, Sharad Malik
September 2004 **Proceedings of the 2nd IEEE/ACM/IFIP international conference on Hardware/software codesign and system synthesis**
Full text available: pdf(300.46 KB) Additional Information: full citation, abstract, references, index terms

In multiprocessor based SoCs, optimizing the communication architecture is often as important as, if not more than, optimizing the computation architecture. While there are mature platforms and techniques for the modeling and evaluation of computation architectures, the same is not true for the communication architectures. A major challenge in modeling the communication architecture is managing the concurrency at multiple levels: at the operation level, multiple communication operations may be a ...

Keywords: bus, design exploration, multiprocessor system, on-chip communication architecture, packet-switching network, retargetable simulation, simulator synthesis

6 Piranha: a scalable architecture based on single-chip multiprocessing
Luiz André Barroso, Kourosh Gharachorloo, Robert McNamara, Andreas Nowatzyk, Shaz Qadeer, Barton Sano, Scott Smith, Robert Stets, Ben Verghese
May 2000 **ACM SIGARCH Computer Architecture News , Proceedings of the 27th annual international symposium on Computer architecture**, Volume 28 Issue 2
Full text available: pdf(191.10 KB) Additional Information: full citation, abstract, references, citings, index terms

The microprocessor industry is currently struggling with higher development costs and longer design times that arise from exceedingly complex processors that are pushing the limits of instruction-level parallelism. Meanwhile, such designs are especially ill suited for important commercial applications, such as on-line transaction processing (OLTP), which suffer from large memory stall times and exhibit little instruction-level parallelism. Given that commercial applications constitute by fa ...

7 Addressing the system-on-a-chip interconnect woes through communication-based design
M. Sgroi, M. Sheets, A. Mihal, K. Keutzer, S. Malik, J. Rabaey, A. Sangiovanni-Vencentelli

June 2001 **Proceedings of the 38th conference on Design automation**

Full text available: pdf(180.03 KB) Additional Information: full citation, abstract, references, citings, index terms

Communication-based design represents a formal method approach to of system-on-a-chip design that considers communication between components as important as the computations they perform. "Our network-on-chip&rdqo ; approach partitions the communication into layers to maximize reuse and provide a programmer with an abstraction of the underlying communication framework. This layered approach is cast in the structure advocated by the OSI Reference network model and is demonstrated with ...

Keywords: communication-based design, network-on-chip, platform-based design, protocol stack

8 Architecture and design of AlphaServer GS320

Kourosh Gharachorloo, Madhu Sharma, Simon Steely, Stephen Van Doren

November 2000 **Proceedings of the ninth international conference on Architectural support for programming languages and operating systems**, Volume 28 , 34 Issue 5 , 5

Full text available: pdf(413.91 KB) Additional Information: full citation, abstract, references, citings, index terms

This paper describes the architecture and implementation of the AlphaServer GS320, a cache-coherent non-uniform memory access multiprocessor developed at Compaq. The AlphaServer GS320 architecture is specifically targeted at medium-scale multiprocessing with 32 to 64 processors. Each node in the design consists of four Alpha 21264 processors, up to 32GB of coherent memory, and an aggressive IO subsystem. The current implementation supports up to 8 such nodes for a total of 32 processors. While s ...

9 Formal specification and design of a message router

Christian Creveuil, Gruia-Catalin Roman

October 1994 **ACM Transactions on Software Engineering and Methodology (TOSEM)**, Volume 3 Issue 4

Full text available: pdf(2.49 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Formal derivation refers to a family of design techniques that entail the development of programs which are guaranteed to be correct by construction. Only limited industrial use of such techniques (e.g., UNITY-style specification refinement) has been reported in the literature, and there is a great need for methodological developments aimed at facilitating their application to complex problems. This article examines the formal specification and design of a message router in an attempt to id ...

Keywords: UNITY, formal methods, program derivation, specification refinement

10 Architecture and design of AlphaServer GS320

Kourosh Gharachorloo, Madhu Sharma, Simon Steely, Stephen Van Doren

November 2000 **ACM SIGPLAN Notices**, Volume 35 Issue 11

Full text available: pdf(1.67 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper describes the architecture and implementation of the AlphaServer GS320, a cache-coherent non-uniform memory access multiprocessor developed at Compaq. The AlphaServer GS320 architecture is specifically targeted at medium-scale multiprocessing with 32 to 64 processors. Each node in the design consists of four Alpha 21264 processors, up to 32GB of coherent memory, and an aggressive IO subsystem. The current

implementation supports up to 8 such nodes for a total of 32 processors. While s ...

11 A hierarchical modeling framework for on-chip communication architectures
Xinping Zhu, Sharad Malik
November 2002 **Proceedings of the 2002 IEEE/ACM international conference on Computer-aided design**
Full text available: pdf(124.52 KB) Additional Information: full citation, abstract, references, citings, index terms

The communication sub-system of complex IC systems is increasingly critical for achieving system performance. Given this, it is important that the on-chip communication architecture should be included in any quantitative evaluation of system design during design space exploration. While there are several mature methodologies for the modeling and evaluation of architectures of processing elements, there is relatively little work done in modeling of an extensive range of on-chip communication arch ...

12 Synchronization and communication in the T3E multiprocessor
Steven L. Scott
September 1996 **Proceedings of the seventh international conference on Architectural support for programming languages and operating systems**, Volume 31 , 30 Issue 9 , 5
Full text available: pdf(1.34 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper describes the synchronization and communication primitives of the Cray T3E multiprocessor, a shared memory system scalable to 2048 processors. We discuss what we have learned from the T3D project (the predecessor to the T3E) and the rationale behind changes made for the T3E. We include performance measurements for various aspects of communication and synchronization.The T3E augments the memory interface of the DEC 21164 microprocessor with a large set of explicitly-managed, external r ...

13 Low power SOCs and NOCs: Plug-in of power models in the StepNP exploration platform: analysis of power/performance trade-offs
Giovanni Beltrame, Gianluca Palermo, Donatella Sciuto, Cristina Silvano
September 2004 **Proceedings of the 2004 international conference on Compilers, architecture, and synthesis for embedded systems**
Full text available: pdf(223.30 KB) Additional Information: full citation, abstract, references, index terms

In this paper, we propose a power/performance estimation layer designed for StepNP, a system-level architecture simulation and exploration platform for Network Processors and Multi-Processor Systems-on-Chip (MP-SoCs). The first goal of our work is to plug-in PIRATE, a parameterizable Network on-Chip in the StepNP platform, to support a fast exploration of on-chip interconnection networks. Up to now, StepNP does not provide any energy profiling, so our second goal is to dynamically plug-in power ...

Keywords: low-power design, multiprocessor, network on chip, platform based design

14 Survey of commercial parallel machines
Gowri Ramanathan, Joel Oren
June 1993 **ACM SIGARCH Computer Architecture News**, Volume 21 Issue 3
Full text available: pdf(1.64 MB) Additional Information: full citation, abstract, citings, index terms

We have presented in this paper the survey of the parallel machines that are marketed today. The survey includes the latest machines available from Kendell Square Research, Thinking Machines Corporation, MasPar Computer Corporation, NCUBE Corporation, Sequent Computer Systems and Parsytec. We have provided the topology, architecture,

cache coherence, synchronization and performance in MFLOPs for each of the machines subject to the availability of information.

15 An overview of the BlueGene/L Supercomputer

NR Adiga, G Almasi, GS Almasi, Y Aridor, R Barik, D Beece, R Bellofatto, G Bhanot, R Bickford, M Blumrich, AA Bright, J Brunheroto, C Caşcaval, J Castaños, W Chan, L Ceze, P Coteus, S Chatterjee, D Chen, G Chiu, TM Cipolla, P Crumley, KM Desai, A Deutsch, T Domany, MB Dombrowa, W Donath, M Eleftheriou, C Erway, J Esch, B Fitch, J Gagliano, A Gara, R Garg, R Germain, ME Giampapa, B Gopalsamy, J Gunnels, M Gupta, F Gustavson, S Hall, RA Haring, D Heidel, P Heidelberger, LM Herger, D Hoenicke, RD Jackson, T Jamal-Eddine, GV Kopcsay, E Krevat, MP Kurhekar, AP Lanzetta, D Lieber, LK Liu, M Lu, M Mendell, A Misra, Y Moatti, L Mok, JE Moreira, BJ Nathanson, M Newton, M Ohmacht, A Oliner, V Pandit, RB Pudota, R Rand, R Regan, B Rubin, A Ruehli, S Rus, RK Sahoo, A Sanomiya, E Schenfeld, M Sharma, E Shmueli, S Singh, P Song, V Srinivasan, BD Steinmacher-Burow, K Strauss, C Surovic, R Swetz, T Takken, RB Tremaine, M Tsao, AR Umamaheshwaran, P Verma, P Vranas, TJC Ward, M Wazlowski, W Barrett, C Engel, B Drehmel, B Hilgart, D Hill, F Kasemkhani, D Krolak, CT Li, T Liebsch, J Marcella, A Muff, A Okomo, M Rouse, A Schram, M Tubbs, G Ulsh, C Wait, J Wittrup, M Bae, K Dockser, L Kissel, MK Seager, JS Vetter, K Yates

November 2002 **Proceedings of the 2002 ACM/IEEE conference on Supercomputing**

Full text available: pdf(357.61 KB) Additional Information: full citation, abstract, references, citings, index terms

This paper gives an overview of the BlueGene/L Supercomputer. This is a jointly funded research partnership between IBM and the Lawrence Livermore National Laboratory as part of the United States Department of Energy ASCI Advanced Architecture Research Program. Application performance and scaling studies have recently been initiated with partners at a number of academic and government institutions, including the San Diego Supercomputer Center and the California Institute of Technology. This mass ...

16 ipChinook: an integrated IP-based design framework for distributed embedded systems

Pai Chou, Ross Ortega, Ken Hines, Kurt Patridge, Gaetano Borriello

June 1999 **Proceedings of the 36th ACM/IEEE conference on Design automation**

Full text available: pdf(89.16 KB) Additional Information: full citation, references, citings, index terms

17 Gilgamesh: a multithreaded processor-in-memory architecture for petaflops computing

Thomas L. Sterling, Hans P. Zima

November 2002 **Proceedings of the 2002 ACM/IEEE conference on Supercomputing**

Full text available: pdf(322.86 KB) Additional Information: full citation, abstract, references, citings, index terms

Processor-in-Memory (PIM) architectures avoid the von Neumann bottleneck in conventional machines by integrating high-density DRAM and CMOS logic on the same chip. Parallel systems based on this new technology are expected to provide higher scalability, adaptability, robustness, fault tolerance and lower power consumption than current MPPs or commodity clusters. In this paper we describe the design of *Gilgamesh*, a PIM-based massively parallel architecture, and elements of its execution mo ...

Keywords: Petaflops computing, Processor-In-Memory, data parallel processing, irregular applications, parallel architectures

18 Tutorial on parallel processing for design automation applications (tutorial session)

J. M. Hancock, S. DasGupta

July 1986 **Proceedings of the 23rd ACM/IEEE conference on Design automation**

Full text available: pdf(927.30 KB) Additional Information: full citation, abstract, references, index terms

This tutorial is designed as an introduction to the field of parallel processing and to its impact on the field of design automation (DA). It starts by reviewing the history of parallel processing. Examples of current hardware are discussed, focusing on the trade offs between custom hardware and general purpose hardware. Interconnection of processors and memory is discussed to demonstrate the range of architectural options available. Then, two key DA algorithms, implemented on parallel mach ...



19 Session summaries from the 17th symposium on operating systems principle (SOSP'99)

Jay Lepreau, Eric Eide

April 2000 **ACM SIGOPS Operating Systems Review**, Volume 34 Issue 2

Full text available: pdf(3.15 MB) Additional Information: full citation, index terms



20 Spinach: a liberty-based simulator for programmable network interface architectures

Paul Willmann, Michael Brogioli, Vijay S. Pai

June 2004 **ACM SIGPLAN Notices , Proceedings of the 2004 ACM SIGPLAN/SIGBED conference on Languages, compilers, and tools for embedded systems**, Volume 39 Issue 7

Full text available: pdf(336.99 KB) Additional Information: full citation, abstract, references, citings, index terms

This paper presents Spinach, a new simulator toolset specifically designed to target programmable network interface architectures. Spinach models both system components that are common to all programmable environments (e.g., ALUs, control and data paths, registers, instruction processing) and components that are specific to the embedded systems and network interface environments (e.g., software-controlled scratchpad memory, hardware assists for DMA and medium access control).Spinach is built on ...

Keywords: embedded systems, programmable network interfaces, simulation

Results 1 - 20 of 52 Result page: **1** 2 3 next

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.
Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player

PORTAL
USPTO
Subscribe (Full Service) Register (Limited Service, Free) Login
Search: The ACM Digital Library The Guide
+packet +order +arbitrating +interconnect +design +router +
SEARCH
Feedback Report a problem Satisfaction survey

Terms used
packet order arbitrating interconnect design router target language hardware

Found **39** of **158,639**

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 39 Result page: **1** 2 next

Relevance scale

1 A comparative study of arbitration algorithms for the Alpha 21364 pipelined router
Shubhendu S. Mukherjee, Federico Silla, Peter Bannon, Joel Emer, Steve Lang, David Webb
October 2002 **Proceedings of the 10th international conference on Architectural support for programming languages and operating systems**, Volume 30 , 36 , 37 Issue 5 , 5 , 10
Full text available: pdf(1.44 MB) Additional Information: full citation, abstract, references

Interconnection networks usually consist of a fabric of interconnected routers, which receive packets arriving at their input ports and forward them to appropriate output ports. Unfortunately, network packets moving through these routers are often delayed due to conflicting demand for resources, such as output ports or buffer space. Hence, routers typically employ *arbiters* that resolve conflicting resource demands to maximize the number of matches between packets waiting at input ports an ...

2 Piranha: a scalable architecture based on single-chip multiprocessing
Luiz André Barroso, Kourosh Gharachorloo, Robert McNamara, Andreas Nowatzyk, Shaz Qadeer, Barton Sano, Scott Smith, Robert Stets, Ben Verghese
May 2000 **ACM SIGARCH Computer Architecture News , Proceedings of the 27th annual international symposium on Computer architecture**, Volume 28 Issue 2
Full text available: pdf(191.10 KB) Additional Information: full citation, abstract, references, citings, index terms

The microprocessor industry is currently struggling with higher development costs and longer design times that arise from exceedingly complex processors that are pushing the limits of instruction-level parallelism. Meanwhile, such designs are especially ill suited for important commercial applications, such as on-line transaction processing (OLTP), which suffer from large memory stall times and exhibit little instruction-level parallelism. Given that commercial applications constitute by fa ...

3 Special session on reconfigurable computing: The happy marriage of architecture and application in next-generation reconfigurable systems
Ingrid Verbauwhede, Patrick Schaumont
April 2004 **Proceedings of the 1st conference on Computing frontiers**
Full text available: pdf(398.28 KB) Additional Information: full citation, abstract, references, index terms

New applications and standards are first conceived only for functional correctness and without concerns for the target architecture. The next challenge is to map them onto an architecture. Embedding such applications in a portable, low-energy context is the art of

molding it onto an energy-efficient target architecture combined with an energy efficient execution. With a reconfigurable architecture, this task becomes a two-way process where the architecture adapts to the application and vice-vers ...

Keywords: embedded, real-time systems

4 Balancing performance and flexibility with hardware support for network architectures
Ilija Hadžić, Jonathan M. Smith
November 2003 **ACM Transactions on Computer Systems (TOCS)**, Volume 21 Issue 4

Full text available: pdf(719.03 KB) Additional Information: full citation, abstract, references, index terms

The goals of performance and flexibility are often at odds in the design of network systems. The tension is common enough to justify an architectural solution, rather than a set of context-specific solutions. The Programmable Protocol Processing Pipeline (P4) design uses programmable hardware to selectively accelerate protocol processing functions. A set of field-programmable gate arrays (FPGAs) and an associated library of network processing modules implemented in hardware are augmented with so ...

Keywords: FPGA, P4, computer networking, flexibility, hardware, performance, programmable logic devices, programmable networks, protocol processing

5 Multiprocessor SoC: design strategies and programming models: Modeling operation and microarchitecture concurrency for communication architectures with application to retargetable simulation
Xinping Zhu, Wei Qin, Sharad Malik
September 2004 **Proceedings of the 2nd IEEE/ACM/IFIP international conference on Hardware/software codesign and system synthesis**

Full text available: pdf(300.46 KB) Additional Information: full citation, abstract, references, index terms

In multiprocessor based SoCs, optimizing the communication architecture is often as important as, if not more than, optimizing the computation architecture. While there are mature platforms and techniques for the modeling and evaluation of computation architectures, the same is not true for the communication architectures. A major challenge in modeling the communication architecture is managing the concurrency at multiple levels: at the operation level, multiple communication operations may be a ...

Keywords: bus, design exploration, multiprocessor system, on-chip communication architecture, packet-switching network, retargetable simulation, simulator synthesis

6 Special issue: AI in engineering
D. Sriram, R. Joobbani
January 1985 **ACM SIGART Bulletin**, Issue 91

Full text available: pdf(8.79 MB) Additional Information: full citation, abstract

The papers in this special issue were compiled from responses to the announcement in the July 1984 issue of the SIGART newsletter and notices posted over the ARPAnet. The interest being shown in this area is reflected in the sixty papers received from over six countries. About half the papers were received over the computer network.

7 Formal specification and design of a message router
Christian Creveuil, Gruia-Catalin Roman
October 1994 **ACM Transactions on Software Engineering and Methodology (TOSEM)**, Volume 3 Issue 4

Full text available: pdf(2.49 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Formal derivation refers to a family of design techniques that entail the development of programs which are guaranteed to be correct by construction. Only limited industrial use of such techniques (e.g., UNITY-style specification refinement) has been reported in the literature, and there is a great need for methodological developments aimed at facilitating their application to complex problems. This article examines the formal specification and design of a message router in an attempt to id ...

Keywords: UNITY, formal methods, program derivation, specification refinement

8 Architecture and design of AlphaServer GS320

Kourosh Gharachorloo, Madhu Sharma, Simon Steely, Stephen Van Doren

November 2000 **Proceedings of the ninth international conference on Architectural support for programming languages and operating systems**, Volume 28 , 34 Issue 5 , 5

Full text available: pdf(413.91 KB) Additional Information: full citation, abstract, references, citings, index terms

This paper describes the architecture and implementation of the AlphaServer GS320, a cache-coherent non-uniform memory access multiprocessor developed at Compaq. The AlphaServer GS320 architecture is specifically targeted at medium-scale multiprocessing with 32 to 64 processors. Each node in the design consists of four Alpha 21264 processors, up to 32GB of coherent memory, and an aggressive IO subsystem. The current implementation supports up to 8 such nodes for a total of 32 processors. While s ...

9 Architecture and design of AlphaServer GS320

Kourosh Gharachorloo, Madhu Sharma, Simon Steely, Stephen Van Doren

November 2000 **ACM SIGPLAN Notices**, Volume 35 Issue 11

Full text available: pdf(1.67 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper describes the architecture and implementation of the AlphaServer GS320, a cache-coherent non-uniform memory access multiprocessor developed at Compaq. The AlphaServer GS320 architecture is specifically targeted at medium-scale multiprocessing with 32 to 64 processors. Each node in the design consists of four Alpha 21264 processors, up to 32GB of coherent memory, and an aggressive IO subsystem. The current implementation supports up to 8 such nodes for a total of 32 processors. While s ...

10 Addressing the system-on-a-chip interconnect woes through communication-based design

M. Sgroi, M. Sheets, A. Mihal, K. Keutzer, S. Malik, J. Rabaey, A. Sangiovanni-Vencentelli

June 2001 **Proceedings of the 38th conference on Design automation**

Full text available: pdf(180.03 KB) Additional Information: full citation, abstract, references, citings, index terms

Communication-based design represents a formal method approach to of system-on-a-chip design that considers communication between components as important as the computations they perform. "Our network-on-chip&rdqo ; approach partitions the communication into layers to maximize reuse and provide a programmer with an abstraction of the underlying communication framework. This layered approach is cast in the structure advocated by the OSI Reference network model and is demonstrated with ...

Keywords: communication-based design, network-on-chip, platform-based design, protocol stack

11 Session summaries from the 17th symposium on operating systems principle (SOSP'99)
Jay Lepreau, Eric Eide
April 2000 **ACM SIGOPS Operating Systems Review**, Volume 34 Issue 2
Full text available: pdf(3.15 MB) Additional Information: full citation, index terms

12 An overview of the BlueGene/L Supercomputer
NR Adiga, G Almasi, GS Almasi, Y Aridor, R Barik, D Beece, R Bellofatto, G Bhanot, R Bickford, M Blumrich, AA Bright, J Brunheroto, C Caşcaval, J Castaños, W Chan, L Ceze, P Coteus, S Chatterjee, D Chen, G Chiu, TM Cipolla, P Crumley, KM Desai, A Deutsch, T Domany, MB Dombrowa, W Donath, M Eleftheriou, C Erway, J Esch, B Fitch, J Gagliano, A Gara, R Garg, R Germain, ME Giampapa, B Gopalsamy, J Gunnels, M Gupta, F Gustavson, S Hall, RA Haring, D Heidel, P Heidelberger, LM Herger, D Hoenicke, RD Jackson, T Jamal-Eddine, GV Kopcsay, E Krevat, MP Kurhekar, AP Lanzetta, D Lieber, LK Liu, M Lu, M Mendell, A Misra, Y Moatti, L Mok, JE Moreira, BJ Nathanson, M Newton, M Ohmacht, A Oliner, V Pandit, RB Pudota, R Rand, R Regan, B Rubin, A Ruehli, S Rus, RK Sahoo, A Sanomiya, E Schenfeld, M Sharma, E Shmueli, S Singh, P Song, V Srinivasan, BD Steinmacher-Burow, K Strauss, C Surovic, R Swetz, T Takken, RB Tremaine, M Tsao, AR Umamaheshwaran, P Verma, P Vranas, TJC Ward, M Wazlowski, W Barrett, C Engel, B Drehmel, B Hilgart, D Hill, F Kasemkhani, D Krolak, CT Li, T Liebsch, J Marcella, A Muff, A Okomo, M Rouse, A Schram, M Tubbs, G Ulsh, C Wait, J Wittrup, M Bae, K Dockser, L Kissel, MK Seager, JS Vetter, K Yates
November 2002 **Proceedings of the 2002 ACM/IEEE conference on Supercomputing**
Full text available: pdf(357.61 KB) Additional Information: full citation, abstract, references, citings, index terms

This paper gives an overview of the BlueGene/L Supercomputer. This is a jointly funded research partnership between IBM and the Lawrence Livermore National Laboratory as part of the United States Department of Energy ASCI Advanced Architecture Research Program. Application performance and scaling studies have recently been initiated with partners at a number of academic and government institutions, including the San Diego Supercomputer Center and the California Institute of Technology. This mass ...

13 Synchronization and communication in the T3E multiprocessor
Steven L. Scott
September 1996 **Proceedings of the seventh international conference on Architectural support for programming languages and operating systems**, Volume 31 , 30 Issue 9 , 5
Full text available: pdf(1.34 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper describes the synchronization and communication primitives of the Cray T3E multiprocessor, a shared memory system scalable to 2048 processors. We discuss what we have learned from the T3D project (the predecessor to the T3E) and the rationale behind changes made for the T3E. We include performance measurements for various aspects of communication and synchronization.The T3E augments the memory interface of the DEC 21164 microprocessor with a large set of explicitly-managed, external r ...

14 Survey of commercial parallel machines
Gowri Ramanathan, Joel Oren
June 1993 **ACM SIGARCH Computer Architecture News**, Volume 21 Issue 3
Full text available: pdf(1.64 MB) Additional Information: full citation, abstract, citings, index terms

We have presented in this paper the survey of the parallel machines that are marketed today. The survey includes the latest machines available from Kendell Square Research,

Thinking Machines Corporation, MasPar Computer Corporation, NCUBE Corporation, Sequent Computer Systems and Parsytec. We have provided the topology, architecture, cache coherence, synchronization and performance in MFLOPs for each of the machines subject to the availability of information.

15 Low power SOCs and NOCs: Plug-in of power models in the StepNP exploration platform: analysis of power/performance trade-offs

Giovanni Beltrame, Gianluca Palermo, Donatella Sciuto, Cristina Silvano

September 2004 **Proceedings of the 2004 international conference on Compilers, architecture, and synthesis for embedded systems**

Full text available: pdf(223.30 KB) Additional Information: full citation, abstract, references, index terms

In this paper, we propose a power/performance estimation layer designed for StepNP, a system-level architecture simulation and exploration platform for Network Processors and Multi-Processor Systems-on-Chip (MP-SoCs). The first goal of our work is to plug-in PIRATE, a parameterizable Network on-Chip in the StepNP platform, to support a fast exploration of on-chip interconnection networks. Up to now, StepNP does not provide any energy profiling, so our second goal is to dynamically plug-in power ...

Keywords: low-power design, multiprocessor, network on chip, platform based design

16 Building blocks for data flow prototypes

Jack B. Dennis, G. Andrew Boughton, Clement K.C. Leung

May 1980 **Proceedings of the 7th annual symposium on Computer Architecture**

Full text available: pdf(614.42 KB) Additional Information: full citation, abstract, references, citings, index terms

A variety of proposed architectures for data flow computers have been advanced. Evaluation of the practical potential of these proposals is being studied through analysis and simulation, but these techniques cannot be used to study a machine design in sufficient detail to make accurate predictions of performance. As a basis for extrapolating cost/performance of these architectures, and for developing a methodology for data flow program preparation, the construction of prototype machines is ...

17 A survey of processors with explicit multithreading

Theo Ungerer, Borut Robič, Jurij Šilc

March 2003 **ACM Computing Surveys (CSUR)**, Volume 35 Issue 1

Full text available: pdf(920.16 KB) Additional Information: full citation, abstract, references, citings, index terms

Hardware multithreading is becoming a generally applied technique in the next generation of microprocessors. Several multithreaded processors are announced by industry or already into production in the areas of high-performance microprocessors, media, and network processors.A multithreaded processor is able to pursue two or more threads of control in parallel within the processor pipeline. The contexts of two or more threads of control are often stored in separate on-chip register sets. Unused i ...

Keywords: Blocked multithreading, interleaved multithreading, simultaneous multithreading

18 Spinach: a liberty-based simulator for programmable network interface architectures

Paul Willmann, Michael Brogioli, Vijay S. Pai

June 2004 **ACM SIGPLAN Notices , Proceedings of the 2004 ACM SIGPLAN/SIGBED conference on Languages, compilers, and tools for embedded systems,** Volume 39 Issue 7

Full text available: pdf(336.99 KB) Additional Information: full citation, abstract, references, citings, index terms

This paper presents Spinach, a new simulator toolset specifically designed to target programmable network interface architectures. Spinach models both system components that are common to all programmable environments (e.g., ALUs, control and data paths, registers, instruction processing) and components that are specific to the embedded systems and network interface environments (e.g., software-controlled scratchpad memory, hardware assists for DMA and medium access control).Spinach is built on ...

Keywords: embedded systems, programmable network interfaces, simulation

19 Gilgamesh: a multithreaded processor-in-memory architecture for petaflops computing

Thomas L. Sterling, Hans P. Zima

November 2002 **Proceedings of the 2002 ACM/IEEE conference on Supercomputing**

Full text available: pdf(322.86 KB) Additional Information: full citation, abstract, references, citings, index terms

Processor-in-Memory (PIM) architectures avoid the von Neumann bottleneck in conventional machines by integrating high-density DRAM and CMOS logic on the same chip. Parallel systems based on this new technology are expected to provide higher scalability, adaptability, robustness, fault tolerance and lower power consumption than current MPPs or commodity clusters. In this paper we describe the design of *Gilgamesh*, a PIM-based massively parallel architecture, and elements of its execution mo ...

Keywords: Petaflops computing, Processor-In-Memory, data parallel processing, irregular applications, parallel architectures

20 Timestamp snooping: an approach for extending SMPs

Milo M. K. Martin, Daniel J. Sorin, Anatassia Ailamaki, Alaa R. Alameldeen, Ross M. Dickson, Carl J. Mauer, Kevin E. Moore, Manoj Plakal, Mark D. Hill, David A. Wood

November 2000 **Proceedings of the ninth international conference on Architectural support for programming languages and operating systems**, Volume 28 , 34 Issue 5 , 5

Full text available: pdf(164.27 KB) Additional Information: full citation, abstract, references, citings, index terms

Symmetric muultiprocessor (SMP) servers provide superior performance for the commercial workloads that dominate the Internet. Our simulation results show that over one-third of cache misses by these applications result in cache-to-cache transfers, where the data is found in another processor's cache rather than in memory. SMPs are optimized for this case by using snooping protocols that broadcast address transactions to all processors. Conversely, directory-based shared-memory systems must indir ...

Results 1 - 20 of 39 Result page: **1** 2 next

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.
Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player